

CONCURRENT SCHEDULES OF POSITIVE AND NEGATIVE REINFORCEMENT:
DIFFERENTIAL-IMPACT AND DIFFERENTIAL-OUTCOMES HYPOTHESES

MICHAEL A. MAGOON

AUBURN UNIVERSITY

AND

THOMAS S. CRITCHFIELD

ILLINOIS STATE UNIVERSITY

Considerable evidence from outside of operant psychology suggests that aversive events exert greater influence over behavior than equal-sized positive-reinforcement events. Operant theory is largely moot on this point, and most operant research is uninformative because of a scaling problem that prevents aversive events and those based on positive reinforcement from being directly compared. In the present investigation, humans' mouse-click responses were maintained on similarly structured, concurrent schedules of positive (money gain) and negative (avoidance of money loss) reinforcement. Because gains and losses were of equal magnitude, according to the analytical conventions of the generalized matching law, bias ($\log b \neq 0$) would indicate differential impact by one type of consequence; however, no systematic bias was observed. Further research is needed to reconcile this outcome with apparently robust findings in other literatures of superior behavior control by aversive events. In an incidental finding, the linear function relating log behavior ratio and log reinforcement ratio was steeper for concurrent negative and positive reinforcement than for control conditions involving concurrent positive reinforcement. This may represent the first empirical confirmation of a free-operant differential-outcomes effect predicted by contingency-discriminability theories of choice.

Key words: positive reinforcement, negative reinforcement, concurrent schedules, differential-outcomes effect, generalized matching law, humans, mouse click, money

A law of effect, in which consequences influence the behavior upon which they are contingent, is inherent in all conceptions of operant behavior. Whether all consequences function similarly in such a law has, however, been much debated. At issue is the validity of a symmetrical law of effect (e.g., Farley & Fantino, 1978; Thorndike, 1911), which assigns parallel (i.e., equal but opposite) functions to positive reinforcement and conse-

quences based on aversive events (punishment and negative reinforcement).

The precepts of a symmetrical law of effect can be challenged on two levels. First, two-factor theories of punishment (e.g., Dinsmoor, 1954) and negative reinforcement (e.g., Mowrer, 1947) invoke a *differential-mechanism* hypothesis in which, compared to positive reinforcement, punishment and negative reinforcement are held to employ different, or additional, behavioral mechanisms. By contrast, one-factor theories (Herrnstein & Hine-line, 1966; Rachlin & Herrnstein, 1969; Sidman, 1962; Thorndike, 1911) endorse a *common-mechanism* hypothesis. That is, appetitive and aversive consequences are viewed as oppositely valenced components of a single behavioral process. For the most part, the debate about differential versus common mechanism remains unresolved (e.g., Critchfield, Paletz, MacAleese, & Newland, 2003; Dinsmoor, 2001; Hine-line, 1984).

Second, some interpretative writings within behavior analysis appear to endorse a *differential-impact* hypothesis asserting that, whatever

This research was part of Mike Magoon's Master's thesis at Auburn University and was supported by a grant-in-aid to Tom Critchfield from the Auburn University College of Liberal Arts. Magoon was partially supported by the Illinois State University Department of Psychology during data analysis and manuscript preparation, for which we are grateful to Chair David Barone. We thank Mei-Sho Jang for capable computer programming; Chris Newland for advice on the development of the study; Sam Catanzaro for alerting us to relevant literature; and Shawn Gilbert, Brian Wynn, Holly Thompson, Matt Alderson, and Rykert Toledano for help with data collection.

Address correspondence to: Tom Critchfield, Department of Psychology, Illinois State University, Normal, IL 61790 (e-mail: tscritic@ilstu.edu).

doi: 10.1901/jeab.2008.90-1

the mechanism of action, appetitive and aversive consequences exert differing degrees of control over behavior. Skinner (1953), for example, suggested that the effects of punishment are weaker than the effects of positive reinforcement. Daniels (1994) asserted that negative reinforcement is less effective than positive reinforcement: "Negative reinforcement generates just enough behavior to escape or avoid punishment...[whereas] positive reinforcement generates more behavior than is minimally required" (p. 28; italics added). An alternative view would be a *common-impact* hypothesis maintaining that, whatever the mechanism of action, appetitive and aversive consequences exert equal degrees of control over behavior (e.g., Farley & Fantino, 1978).

Few published studies have directly assessed the relative impact on behavior of positive reinforcement versus consequences based on aversive events. Some reports describe effects that may support a differential-impact hypothesis, albeit one that runs counter to Skinner's (1953) ideas on the matter. Specifically, in several studies children's discrimination learning apparently was better promoted by punishing errors than by reinforcing correct responses (e.g., Meyer & Offenbach, 1962; Penney, 1968; Penney & Lupton, 1961; Spence, 1966; Tindall & Ratliff, 1974). Yet qualitatively dissimilar events formed the basis of punishment (e.g., a loud noise) and reinforcement (e.g., food), making comparison on a unit-by-unit basis impossible (e.g., see Farley & Fantino, 1978). In a very few cases, similar effects have arisen when the consequences were of equal nominal value. For example, in a study involving children, Costantini and Hoving (1973) found that acquisition of performance under a differential reinforcement of response duration contingency was better promoted by removing tokens for fast performances than by providing tokens for slow performances. The procedure was brief (five trials, lasting a few minutes), so the durability of the effects is unknown, and to our knowledge the effect has not been replicated with other tasks and types of subjects.

To the extent that punishment can be considered an "unpleasant event," reports from a variety of research areas echo the effects just described. Compared to pleasant events, unpleasant events have tended to evoke relatively more attention; stronger and longer-

lasting changes in mood and emotion; more potent electrophysiological responses; better recognition memory; and more, and better elaborated, causal attributions (for reviews, see Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001; Rozin & Royzman, 2001; Taylor, 1991). It should be noted, however, that the relevance of these pleasant and unpleasant events to operant consequences remains undemonstrated, and the pleasant and unpleasant events under consideration were qualitatively dissimilar, making their comparison on a unit-by-unit basis impossible.

The clearest evidence for a differential-impact effect comes from cognitive-decision studies, which typically present subjects with choices between hypothetical outcomes that bear conceptual similarity to operant consequences (e.g., Rachlin, 1989; Rachlin, Logue, Gibbon, & Frankel, 1986). These studies have been interpreted as indicating that the prospect of a loss weighs more heavily upon decision than the prospect of a gain (for a seminal statement of this conclusion, see Kahneman & Tversky, 1979). Numerous everyday phenomena support laboratory findings in this regard. For example, for purposes of settling a bet, most people judge the equal odds of winning and losing associated with a coin toss as unfavorable (Tversky & Kahneman, 1981), and people typically name a higher price to sell an item they already own (i.e., to "lose" it) than they would be willing to pay to acquire an identical item. Overall, in many decision contexts, a given outcome has greater functional value in a "lose" scenario than in a "gain" scenario (Kahneman, Knetsch, & Thaler, 1990). Importantly, in many relevant studies hypothetical gains and losses were presented in exactly the same scale of measurement (e.g., money amounts), making differential impact easy to assess.

Although no definitive study exists, the preponderance of evidence converges on the notion that, in terms of general psychological potency, "bad [events are] stronger than good [events] in a disappointingly relentless pattern" (Baumeister et al., 2001, p. 362). This conclusion runs counter to a symmetrical law of effect and is opposite to Skinner's (1938) brand of differential-impact hypothesis. To properly evaluate differential-impact hypotheses in an operant context, additional research is needed in which equal amounts of appetitive and aversive consequences (independent var-

iables) are programmed identically to determine whether they exert equal amounts of influence over behavior (dependent variable). Human operant procedures offer a potential advantage in this regard because they typically employ conditioned reinforcers such as money as the basis for consequences (Pilgrim, 1988). Money gains and money losses are measured, quite literally, in the same currency, and thus are readily compared on a unit-by-unit basis.

Concurrent schedules of reinforcement provide a parsimonious way to compare the relative control exerted by two types of consequences (e.g., Herrnstein, 1970). When concurrent-interval schedules produce identical reinforcers, a positive, linear relationship exists between logarithmically transformed response-allocation and reinforcer-availability ratios. This relationship is described by the generalized matching law (Baum, 1974):

$$\log\left(\frac{B_x}{B_y}\right) = a \log\left(\frac{R_x}{R_y}\right) + \log b, \quad (1)$$

in which x and y are the two concurrent response options; B and R , respectively, are measures of behavior allocation and reinforcer availability; and a (slope) and $\log b$ (intercept) are fitted parameters. When the two response options produce different types of reinforcers, relative control by (or preference for) one type is indicated by nonzero values of $\log b$, corresponding to a vertical shift in the linear function, normally described as bias (Baum, 1974, 1979; Chao, 1984; Herrnstein, 1970; McLean & Blampied, 2001; Matthews & Temple, 1979; Miller, 1976).

A direct evaluation of the differential-impact hypothesis might be accomplished using concurrent schedules of positive and negative reinforcement in which the magnitude of the two types of reinforcers is equated. In the present context, negative reinforcement is easier to consider than punishment, which must be superimposed on behavior that is maintained by some other operation (such as positive reinforcement), thereby making pure effects difficult to estimate. In principle, negative reinforcement can operate, and be evaluated, independently.

Equation 1 has been applied to negative reinforcement by considering each cancellation of an aversive event (such as shocks for nonhumans) as a reinforcer, as de Villiers

(1972, 1974) and others have suggested. Evaluated in this way, matching under concurrent schedules of negative reinforcement follows the positively sloped, linear pattern seen under concurrent schedules of positive reinforcement (Baum, 1973; Ferrari & Todorov, 1980; Hutton, Gardner, & Lewis, 1978; Logue & de Villiers, 1978; Poling, 1978). This finding does not bear directly on differential-impact hypotheses, however, because negative and positive reinforcement were not compared via the same matching function. Studies of matching under concurrent schedules of positive versus negative reinforcement also have yielded unremarkable matching functions, but studies involving nonhuman subjects (e.g., Logue & de Villiers, 1981) employed qualitatively dissimilar aversive (shock) and appetitive (food) events, precluding simple comparison on a unit-by-unit basis.

At least two investigations have examined concurrent schedules of positive versus negative reinforcement in humans using equal-sized money outcomes to establish the two types of reinforcers (Ruddle, Bradshaw, & Szabadi, 1981; Ruddle, Bradshaw, Szabadi, & Foster, 1982). Of nine response-matching functions generated by 6 participants, four reflected a bias for negative reinforcement, and five reflected a bias for positive reinforcement. In only two cases, however, did bias differ significantly from zero (one bias for positive reinforcement and one bias for negative reinforcement). Time-matching functions were similar. This mixed evidence was complicated by the fact that reinforcer type was perfectly confounded with operandum location, making it possible that only side bias was observed. Additionally, neither study by Ruddle and colleagues incorporated a control condition involving matching under homogeneous (e.g., all positive or all negative) reinforcement which, as will be illustrated later, may be useful in disentangling different sources of bias.

The present study sought to generate new data relevant to differential-impact hypotheses about positive and negative reinforcement. Participants worked on concurrently programmed, similarly structured, variable-cycle schedules of positive and negative reinforcement. Each subject contributed two complete matching functions, one involving homogeneous (all positive) reinforcement, and one

involving heterogeneous (negative vs. positive) reinforcement, with care taken not to confound reinforcer type with operandum location. Of primary interest was whether systematic bias would be observed.

METHOD

Subjects, Setting, and Apparatus

Four female undergraduates volunteered after reading posted notices seeking participants in research on "choice and decision making." They were accepted into the study upon providing informed consent and on the basis of their ability to attend at least three laboratory visits per week at times when an experimenter was available. The informed consent agreement indicated that money earnings in the experiment depended upon performance and that the investigation was expected to require up to 45 total hours to complete. Earnings accrued through experimental procedures averaged approximately \$5.50 (U.S.) per hour of participation. When the data were collected, the U.S. minimum wage was \$5.15. Each participant received her total earnings after discharge from the study.

One volunteer was removed from the investigation because of erratic attendance and performance in baseline conditions; no data are reported for this individual. One participant (S447) withdrew prematurely, citing boredom with the experimental task, leaving behind a limited series of conditions that are described here. The remaining participants (S448 and S449) completed as many conditions as permitted by their speed in achieving asymptotic performance and by their respective schedules during the academic term during which they worked. To readers who are accustomed to working with nonhuman subjects, the resulting durations of participation (Appendix A) may seem brief; it should be noted, however, that participant retention is a recurring challenge in human operant research (e.g. Pilgrim, 1998), and most human operant experiments are considerably briefer than the present one.

Experimental sessions were conducted in three 3-m by 4-m rooms, each equipped with a table, chair, color computer monitor, and mouse. To mask external noise, soft instrumental music played continuously during experimental sessions from stereo speakers affixed to the

ceiling of the rooms. IBM®-compatible computers in an adjacent room controlled experimental events and collected the data according to a custom program written in the BASIC programming language using QuickBasic®.

Procedure

Participants visited the laboratory 4 or 5 days per week. Experimental sessions lasted 10 min and were separated by short breaks of about 2 to 5 min during which an experimenter recorded the data and initiated the next session. Participants completed about eight sessions during each 2-hr visit to the laboratory.

Experimental task. The concurrent-schedules procedure was based on that of Madden and Perone (1999, 2003; see also Critchfield & Magoon, 2001; Critchfield et al., 2003). Sessions began with display of the message "Click Here To Begin" displayed near the center of the screen directly above a rectangular box with the word "Ready" inside of it. Clicking the box cleared the prompts and produced two rectangles, or work areas, each approximately 13 cm wide by 17 cm high, and each occupying one side of the screen. An arrow-shaped cursor indicated the virtual position of the mouse. Within each of the two work areas was a small (approximately 1.6 square cm) colored target. Throughout the session, targets moved in random directions at a rate of about 1.25 cm/s. Clicks within the borders of a target registered responses upon which the reinforcement schedules were based. Clicks elsewhere were ineffective and were not counted.

At the start of a session, both work areas featured a white background. The side on which a subject's first response occurred remained white and the message "Mouse On" appeared just below it. The background of the other side turned black and the message "Mouse Off" appeared below it. Reinforcement schedules associated with each side continued to operate. Responses on the "Mouse On" side were recorded and entered into the reinforcement contingency that was programmed for that side, whereas responses on the "Mouse Off" side were not recorded and had no programmed effect.

Changeover cost. We were concerned about using a changeover delay (c.f., Ruddle et al., 1982) to discourage adventitious reinforcement of switching between schedules. Under negative reinforcement, the common practice

Table 1
Reinforcement schedule values used in the experiment.

Reinforcement Ratio	Variable-cycle schedule value (s)		Programmed reinforcers per hr	
x:y	x	y	x	y
9:1	11	100	327	36
6:1	12	70	300	51
4:1	13	50	277	72
2:1	15	30	240	120
3:2	17	25	218	144
1:1	20	20	180	180

Note. Designations *x* and *y* refer to the concurrent schedules indexed in Equation 1.

of suspending reinforcement-schedule timers for a period just after a switch between schedules would create a safety period during which no money losses could be experienced, thereby possibly reinforcing changeovers. A fixed-ratio changeover requirement was programmed instead. Located between the two work areas was a 2.5-cm square changeover (CO) button labeled with the word "Change." Clicking the CO button five consecutive times (which pilot work showed could be accomplished in less than 1 s) reversed the status of the two work areas. The formerly inactive screen side became white, and the message "Mouse On" appeared under it. The formerly active screen side became black, and the message "Mouse Off" appeared under it. A click on the active target prior to completing five consecutive changeover responses reset the CO counter to zero. Session and reinforcement timers continued to operate while CO responding took place.

Schedules and consequences. Clicks on the moving targets influenced money gains (in positive-reinforcement schedules) or losses (in negative-reinforcement schedules) according to independent, concurrent variable-cycle schedules (using constant-probability distributions of intervals (Fleshler & Hoffman, 1962)). Cycle schedules function similarly to interval schedules in that reinforcement is contingent upon a single response occurring at a specific point in time (Baron, 1991). In interval schedules, reinforcement is contingent on the first response that occurs after a given time period elapses; responses during the preceding time period are ineffective. In cycle schedules, reinforcement is contingent on the first response that occurs during a time period, with subsequent responses during the time period ineffective (de Villiers, 1972).

Under positive-reinforcement schedules, the first response within a programmed interval immediately caused a money gain to be signaled via a 1.5-s flashing alternation of the most recently clicked target and black text, in the same location, stating "+1.5 ¢" (0.25-s per flash). If no response was made during the interval, the money gain programmed for that interval was cancelled without stimulus change. Responses during the remainder of the interval were ineffective. Under negative-reinforcement schedules, the first response within an interval immediately cancelled the programmed money loss for that interval without stimulus change. If no response was made within a programmed interval, a money loss was signaled at the end of the interval via a 1.5-s flashing alternation of the most recently clicked target and red text, in the same location, stating "-1.5 ¢" (0.25-s per flash). During the money-gain and money-loss messages, the cursor disappeared from the screen and all timers relevant to the procedure were suspended. Cumulative session earnings were not displayed on the screen during the session.

Table 1 shows the schedules that were employed. Subjects completed conditions encompassing a variety of positive- and negative-reinforcement ratios, facilitating data analysis using Equation 1. Appendix A shows the specific schedules (and ratios) experienced by each participant and the sequence in which the conditions were completed. Within a condition, reinforcement schedules made available approximately 360 total reinforcers per hour of session time, aggregated across the two work areas. In homogeneous-reinforcement conditions, positive reinforcement schedules operated in both work areas. In heterogeneous-reinforcement conditions, positive reinforcement operated in one work area,

and negative reinforcement operated in the other.

As Appendix A shows, the design can be thought of as incorporating pairs of conditions, one homogeneous and one heterogeneous, with identical programmed reinforcement ratios. Across heterogeneous-reinforcement conditions, each participant experienced negative reinforcement at least once in each of the two work areas, and at least once as the richer and leaner of the two schedules operating within a condition. In this way consistent relationships were avoided among the work area (left side vs. right side), reinforcement frequency (rich vs. lean schedule), and reinforcement type (positive vs. negative).

Pilot work suggested that participants would not collect all scheduled money gains on positive-reinforcement schedules, or avoid all scheduled money losses on negative-reinforcement schedules. Consequently, net earnings potentially could be low, or even negative, for laboratory visits involving heterogeneous reinforcement. Because of concerns that participants might withdraw from the experiment if net earnings failed to increase across many sessions, money totals produced during a session were supplemented for each heterogeneous-reinforcement session. For consistency across participants, the supplement equaled the programmed session rate of money loss (see Appendix A). The supplements had no bearing on response-reinforcer relations or feedback messages that occurred during experimental sessions, and pilot work suggested that this procedure of supplementing earnings did not systematically affect within-session performance (Critchfield & Magoon, 2001). Although participants were not queried directly, we believe that it is unlikely that they were aware of the supplements. They were not informed that the supplements would be provided, and they asked about their aggregate earning totals only occasionally.

Discriminative stimuli. Within a condition, the schedules programmed in each work area remained constant, and thus were associated with a distinct location. Additionally, the moving target on each side was displayed in a distinctive color, which remained constant within a condition. Color pairs changed across conditions (see Appendix A), and were assigned randomly from a pool of 16 different hues, with the constraint that the same pair of colors could not be used in adjacent conditions.

Instructions and preliminary training. Prior to the start of the main experiment, participants completed approximately 2.5 hr of preliminary training during which they read printed instructions describing the experimental task and received their first exposure to examples of positive- and negative-reinforcement contingencies. Appendix B describes the training and reproduces the instructions.

Stability criteria. Normally, a condition was terminated when one of the following criteria was met: (a) for both response-allocation and time-allocation proportions, over four consecutive sessions, the difference in means between the first and second pair of sessions differed by no more than 10% of the four-session mean; or (b) all response and time proportions in four consecutive sessions were less than 0.1 or greater than 0.9, suggesting floor or ceiling effects. In the event that neither criterion was met within 10 sessions, performance was judged to be stable when visual inspection of graphed data revealed no systematic trend in either response or time allocation proportions.

RESULTS

Appendix A shows the number of sessions required for each participant to complete each of the experimental conditions. Stability was achieved under both the homogeneous-reinforcement schedules (for all conditions for 3 participants, $M = 5.7$ sessions, range = 4 to 12) and heterogeneous-reinforcement schedules ($M = 6.2$ sessions, range = 4 to 11). Terminal data (means from the final four sessions per condition), shown in Appendix A, were used in all analyses. Consistent with previous investigations (e.g., de Villiers, 1972, 1974; Logue & de Villiers, 1978), in the case of negative-reinforcement schedules, a reinforcer was defined as canceling a scheduled money loss.

Analytical strategy. Of primary interest was whether the log b parameter of Equation 1 was different under conditions of homogeneous versus heterogeneous reinforcement. For this purpose we employed an F -test based on the analysis of covariance that is widely used in the biological sciences to compare fitted parameters from two or more data sets (see Motulsky & Christopoulos, 2006; Zar, 1999). The F -test approach avoids some dubious assumptions (see Zar, 1999) that underpin the modified t -

test approach that has been recommended for behavioral data (Davison & McCarthy, 1988).

Comparing the y -intercepts of two linear functions is complicated by the fact that y -intercept covaries with slope (Motulsky & Christopoulos, 2006; Zar, 1999). Thus the F -test approach begins by comparing slopes. For each data set, regressions are performed, yielding three sums of squares (SS ; one each for the x and y variables and one for their product, or interaction), with degrees of freedom (DF) for each equal to the number of observations minus 2. A *common* regression employs the sum of these quantities, producing a common residual SS and a common DF [total number of observations minus [number of data sets minus 1]]. Next, a *pooled* regression is performed on the sum of residual SS of the individual data sets, with DF equal to the sum of residual DF for the individual data sets (number of observations in each set minus 2). The F -ratio, reported with DF (common, pooled), is calculated as follows

$$F = \frac{\left(\frac{SS_{common} - SS_{pooled}}{k - 1} \right)}{\frac{SS_{pooled}}{DF_{pooled}}} \quad (2)$$

with k = number of data sets. Thus, in a two-function analysis like the present one, the numerator reduces to the difference of common and pooled SS .

If the best-fit slopes of two data sets are not significantly different, then an F -test may be used to compare y -intercepts derived from those data sets. If, however, the slopes are significantly different, then y -intercept comparisons are problematic. As will be described, slopes *did* differ for homogeneous and heterogeneous reinforcement functions, thereby precluding direct comparisons of y -intercepts. As an alternative, the y -intercept of each empirical homogeneous-reinforcement function was compared, using the F -test procedure, with a hypothetical function of identical slope and y -intercept equal to that of the corresponding heterogeneous-reinforcement function (this approach is consistent with previous recommendations for comparing empirical and theoretical functions; see Davison & McCarthy, 1988). Because the slopes are identical, it is possible to proceed to the test of intercepts described below; in such cases we

do not report the results of the slope test because $F = 0$.

The analysis of intercepts employs residual SS for a total regression based on an aggregate of all observations (DF = number of total observations minus 2). The resulting F -ratio, reported with DF (total, common), is

$$F = \frac{\left(\frac{SS_{total} - SS_{common}}{k - 1} \right)}{\frac{SS_{common}}{DF_{common}}} \quad (3)$$

Again, as only two data sets are involved, the numerator becomes the difference of the two SS s.

Side bias. A preliminary analysis was conducted to identify any side biases that might occur independently of variations in reinforcer type. Equation 1 was fitted separately to the homogeneous-reinforcement and heterogeneous-reinforcement data of Appendix A using least-squares linear regression, with reinforcement and behavior ratios expressed as left-side schedule/right-side schedule. Side bias unrelated to reinforcer type would be evident in $\log b$ estimates for the homogeneous-reinforcement functions.

Figure 1 summarizes the results of this analysis. Equation 1 accounted for between 78% and 99% of the variance in behavior allocation. Under homogeneous reinforcement (open data points), S447 and S449 showed small left-side biases (positive $\log b$ values), and S448 showed a small right-side bias (negative $\log b$ values). An F -test comparing each empirical function with a hypothetical function of equal slope and $\log b = 0$ revealed that none of these estimates were significantly different from 0 (Table 2, top).

Because this analysis did not focus on reinforcer type per se, Figure 1 cannot be used to evaluate the differential-impact hypothesis. The figure does show, however, that slopes of the heterogeneous-reinforcement functions (filled data points) were significantly elevated compared to those of homogeneous-reinforcement functions (open data points; see Table 3, top). To our knowledge, this effect has no exact precedent because no previous study involving concurrent positive and negative reinforcement has included an all positive-reinforcement comparison function (e.g., Baum, 1973; Ferrari & Todorov,

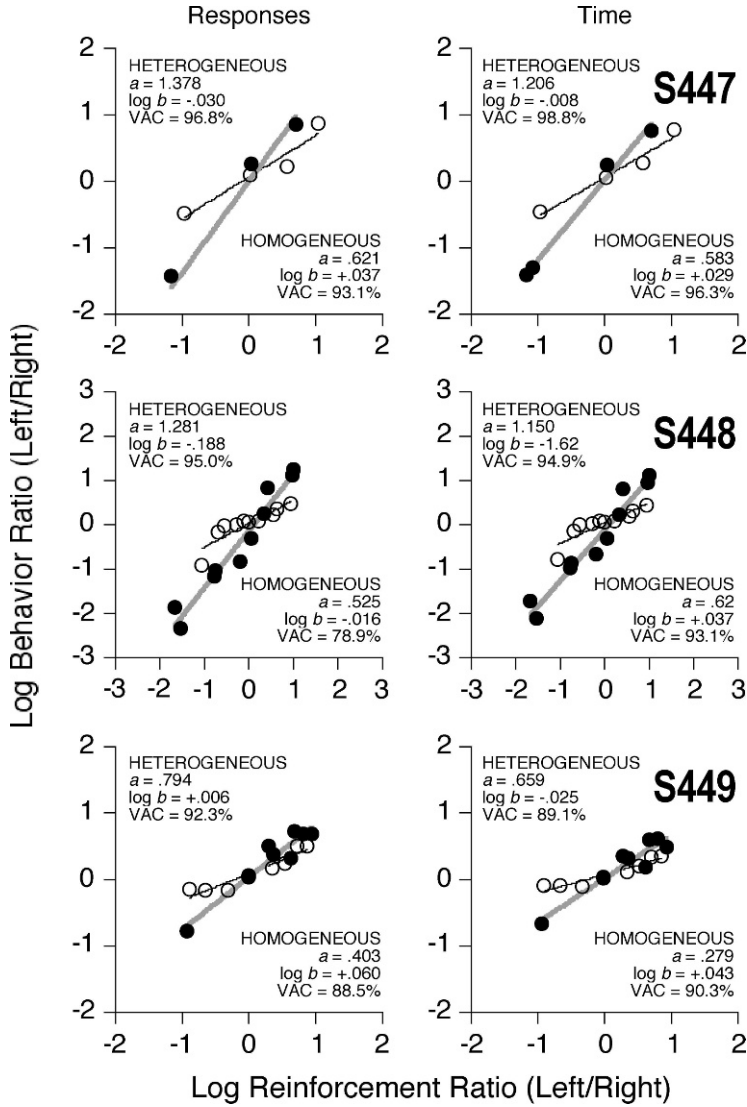


Fig. 1. Allocation of responses (left column) and time (right column) as a function of relative reinforcement frequency for homogeneous-reinforcement (open data points) and heterogeneous-reinforcement (filled data points) conditions. Behavior and reinforcer ratios were organized as left-side schedule/right-side schedule (see text for details). Lines of best fit (thin, black line for homogeneous reinforcement; thick, grey line for heterogeneous reinforcement) and equations describing them were derived using Equation 1 and least-squares linear regression. Note that axes are scaled differently for different participants.

1980; Logue & de Villiers, 1978; Poling, 1978; Ruddle et al., 1982). Although the slope effect does not appear to bear on differential-impact hypotheses, it has other implications regarding theories of operant choice that receive separate attention in the Discussion.

Reinforcer-type bias. As an initial step in evaluating the differential-impact hypothesis, Equation 1 was fitted to the heterogeneous-

reinforcement data in Appendix A, with reinforcement and behavior ratios expressed as negative reinforcement/positive reinforcement, yielding the functions shown in Figure 2. Equation 1 accounted for between 94% and 99% of the variance in behavior allocation. In all cases, Table 3 (bottom) shows that slopes were steeper than for homogeneous reinforcement as per Figure 1 and Table 3

Table 2

Results of comparing empirical bias ($\log b$) estimates with hypothetical functions of equal slope with $\log b = 0$.

Participant	Responses			Time		
	<i>F</i>	<i>DF</i>	<i>p</i>	<i>F</i>	<i>DF</i>	<i>p</i>
<i>Homogeneous: Left/Right (Equation 1)</i>						
S447	0.21	1,5	.662	0.29	1,5	.612
S448	0.81	1,17	.779	<0.01	1,17	.966
S449	1.49	1,13	.243	1.68	1,13	.217
<i>Heterogeneous: Negative/Positive (Equation 1)</i>						
S447	0.62	1,5	.459	0.02	1,5	.892
S448	2.27	1,17	.150	1.77	1,17	.201
S449	0.35	1,13	.560	1.92	1,13	.189
<i>Heterogeneous: Negative/Positive (Equation 4)</i>						
S447	<0.01	1,5	.955	0.52	1,5	.504
S448	0.88	1,17	.360	0.64	1,17	.434
S449	0.05	1,13	.823	0.05	1,13	.502

Note. Top: Estimates derived from fitting Equation 1 to homogeneous-reinforcement data organized as left-side schedule/right-side schedule. See Figure 1 for parameter estimates. Middle: Estimates derived from fitting Equation 1 to heterogeneous-reinforcement data organized as negative-reinforcement schedule/positive-reinforcement schedule. See Figure 2 for parameter estimates. Bottom: Estimates derived from fitting Equation 4 to heterogeneous-reinforcement data organized as negative-reinforcement schedule/positive-reinforcement schedule. See Table 6 for empirical parameter estimates.

(top). Log b estimates were positive (range = +0.01 to +0.19), suggesting that negative reinforcement controlled behavior somewhat more strongly than did positive reinforcement. An F -test comparing each empirical function with a hypothetical function of equal slope and $\log b = 0$ revealed that none of these empirical log b estimates was significantly different from 0 (Table 2, middle). It should be noted, however, that the log b parameter estimates in Figure 2 are impure because the influence of reinforcer type is necessarily intermingled with side-bias effects. That is, for heterogeneous-reinforcement functions, the negative-reinforcement schedule operated sometimes on the left side and sometimes on the right side (see Appendix A), and at the descriptive level small side biases (nonzero log b estimates) were observed for all participants.

Correction for side bias. In recognition of this problem, Equation 1 was fitted to the heterogeneous-reinforcement data in Appendix A, as per the conventions of Figure 2, but with a correction for side bias. Specifically, within a participant's heterogeneous-reinforcement function and prior to regression analysis, the log behavior ratio for each condition was increased or decreased by the log b value from the corresponding homogeneous-reinforcement function in Figure 1. The correction for each condition depended on whether the

negative-reinforcement schedule operated on the left or right screen location. If the former, then the log b estimate from Figure 1 was subtracted from that condition's behavior ratios. If the latter, then the log b estimate from Figure 1 was added to that condition's behavior ratios. As Table 4 shows, all six modified log b estimates were positive (range = +.01 to +.19), which suggests that estimates shown in Figure 2 were not strictly an artifact of side biases. Of these six modified log b estimates, however, only three were significantly different from 0 (Table 4).

DISCUSSION

A Differential-Impact Effect?

At a descriptive level, log b estimates (Figures 1 and 2, Table 4) were consistent with the differential-impact hypothesis that negative reinforcement controls behavior more strongly than positive reinforcement (i.e., "bad is stronger than good;" Baumeister et al., 2001). This outcome should be considered with caution, however, because log b estimates under heterogeneous reinforcement were not always significantly different from zero (Table 4) nor remarkable in an actuarial context. In the latter case, a useful frame of reference is offered by Robinson (1992), who reviewed 34 concurrent-schedule studies, most

Table 3
Results of comparing empirical sensitivity (*a*) estimates for homogeneous-reinforcement and heterogeneous-reinforcement functions.

Participant	Responses			Time		
	<i>F</i>	<i>DF</i>	<i>p</i>	<i>F</i>	<i>DF</i>	<i>p</i>
<i>Left/Right</i>						
S447	12.33	1,4	.025	24.53	1,4	.008
S448	23.85	1,16	<.001	25.46	1,16	<.001
S449	12.81	1,12	.004	14.95	1,12	.002
<i>Negative/Positive</i>						
S447	16.59	1,4	.015	26.58	1,4	.007
S448	29.69	1,16	<.001	30.05	1,16	<.001
S449	15.62	1,12	.002	5.97	1,12	.031

Note. Top: Estimates derived from fitting Equation 1 to both data sets organized as left-side schedule/right-side schedule. See Figure 1 for parameter estimates. Bottom: Estimates derived from fitting Equation1 to homogeneous-reinforcement data organized as left-side schedule/right-side schedule, and heterogeneous-reinforcement data organized as negative-reinforcement schedule/positive-reinforcement schedule. See Figures 1 and 2 for parameter estimates.

of which employed homogeneous reinforcement and therefore would not be expected to show pronounced log *b* estimates. Approximately 60% of individual log *b* values fell into the range of -.10 to +.10, which encompasses most of the log *b* estimates from heterogeneous-reinforcement functions of the present study. Note, too, that in studies involving heterogeneous reinforcement, log *b* estimates can be much larger than those reported here. For instance, in two studies in which the difference in reinforcer type or duration presumably was quite pronounced, absolute log *b* values ranged from 0.27 to 0.81 (food versus electrical brain stimulation in Hollard & Davison, 1971; 6-s versus 2-s food presentations in McLean & Blampied, 2001).

Thus, the present study appears to agree with the findings of Ruddle et al. (1981) and Ruddle et al. (1982) in showing no systematic evidence for a differential-impact effect in concurrent schedules of positive and negative reinforcement. Given the breadth of evidence suggesting that “bad” can be psychologically more potent than “good” (e.g., Baumeister et al., 2001; Kahneman & Tversky, 1979), the question of why no difference was observed deserves some attention. Below we identify four broad issues worth considering.

Does bias always result from heterogeneous reinforcement? The present test of the differential-impact hypothesis rests on the widely held assumption that heterogeneous reinforcement should create bias in matching. Davison and McCarthy (1988) observed that

surprisingly few studies exist to bolster this assumption. In most relevant cases, matching under heterogeneous reinforcement was compared to a theoretical ideal (log *b* = 0) rather than an empirical function derived from homogeneous reinforcement (see Baum, 1974; Chao, 1984; Herrnstein, 1970; Matthews & Temple, 1979; Miller, 1976; Ruddle et al., 1981; Ruddle et al., 1982), and only rarely have efforts been undertaken to determine whether the reported effects are statistically significant, leaving the reported log *b* effects impossible to properly evaluate. Additional studies involving appropriate control conditions are needed to evaluate the validity of this core assumption.

What counts as “bad”? The present test employed free-operant avoidance schedules as a means of establishing the “bad” in a version of the “bad versus good” comparison that is inherent in differential-impact hypotheses. Debate could be mounted regarding whether avoidance contingencies are a proper assay for evaluating differential-impact effects. Although behavior analysts often lump negative reinforcement and punishment together by virtue of their shared reliance on aversive establishing operations (e.g., Sidman, 1989; Skinner, 1953), in previous studies that have been cited in support of differential impact, “bad” events were arranged more similarly to punishment (e.g., Costantini & Hoving, 1973; Kahneman & Tversky, 1979). Perhaps, in the present context, positive and negative reinforcement both qualify simply as reinforcement (“good” events). If so, then, in operant

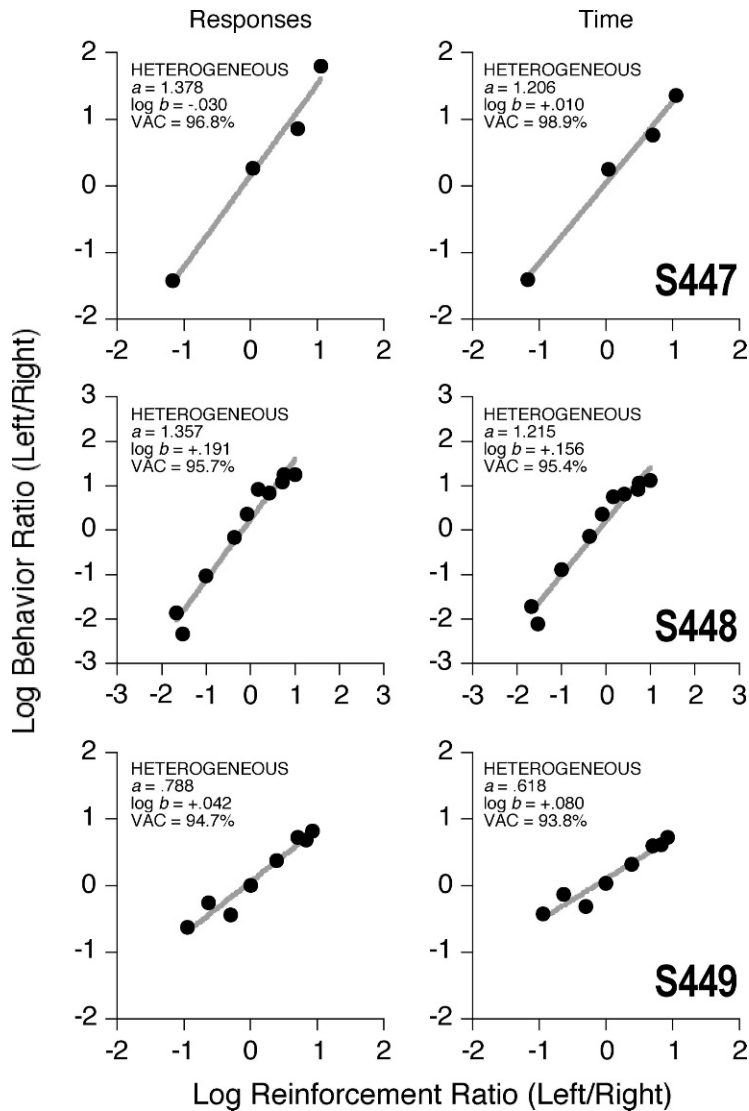


Fig. 2. Allocation of responses (left column) and time (right column) as a function of relative reinforcement frequency, for the heterogeneous reinforcement condition, with ratios organized as negative-reinforcement schedule/positive-reinforcement schedule; see text for details. Lines of best fit and equations describing them were derived using Equation 1 and least-squares linear regression. Note that axes are scaled differently for different participants.

terms, differential-impact hypotheses would invoke only the relative impact on behavior of reinforcement versus punishment. Unfortunately, because punishment cannot be arranged independently of motivating operations such as reinforcement, its impact on behavior is difficult to evaluate in pure form (e.g., Critchfield et al., 2003; Farley & Fantino, 1978).

Reliability of differential-impact effects. It remains possible as well that the disproportion-

ate influence of “bad” events reported in other literatures is a statistical reality that masks substantial individual differences in sensitivity to different kinds of consequences. Many cognitive-decision tasks (e.g., Kahneman & Tversky, 1979) present participants with single choices between hypothetical outcomes, and data are summarized in terms of the percentage of participants who preferred a given alternative. Not all participants show superior control by aversive events (e.g.,

Table 4

Results of fitting Equation 1 to heterogeneous-reinforcement data, with a correction for side bias, organized as negative-reinforcement schedule/positive-reinforcement schedule.

Participant	Behavior measure	Fitted parameters			Log <i>b</i> vs. 0		
		<i>a</i>	log <i>b</i>	%VAC	<i>F</i>	<i>DF</i>	<i>P</i>
S447	Responses	1.390	+0.094	97.6%	0.65	1,5	.456
	Time	1.208	+0.010	99.0%	0.20	1,5	.892
S448	Responses	1.353	+0.187	95.6%	4.64	1,17	.046
	Time	1.214	+0.155	95.5%	3.54	1,17	.077
S449	Responses	0.734	+0.063	94.0%	1.75	1,13	.209
	Time	0.580	+0.095	92.9%	5.32	1,13	.038

Note. Also shown is a comparison of empirical bias (log *b*) estimates with hypothetical functions of equal slope with log *b* = 0. See text for details of the side-bias correction. %VAC = percentage of variance in behavior allocation accounted for by Equation 1.

Kahneman & Tversky, 1979). Perhaps the mixed findings of Ruddle et al. (1981) and Ruddle et al. (1982) accurately represent normal intersubject variation that the present study also would have revealed with more participants. The operant literature does contain hints of individual differences in responsiveness to aversive contingencies. For instance, some nonhuman subjects apparently have difficulty acquiring a shock-avoidance repertoire (e.g., Sidman, 1966). Additionally, in attempting to scale statistically the functional impact of food reinforcers and shock punishers in pigeons, Farley and Fantino (1978) found different relative values for different subjects. Such intersubject differences might be especially pronounced for conditioned consequences such as money, which acquire their capacity to influence behavior through experience that, in the world outside the laboratory, varies across individuals (e.g., Critchfield et al., 2003; Lerman & Vorndran, 2002). The extent to which individual humans differ in their sensitivity to aversive operant consequences is something that only replication with large numbers of participants can reveal.

Difficulties in comparing across literatures. The preceding discussion suggests that comparisons of operant-choice and cognitive-decision studies are complicated by the many procedural differences between the two research traditions (e.g., Christensen, Parker, Silberberg, & Hursh, 1998; Rachlin, 1989). Cognitive-decision studies tend to employ brief, discrete-trials procedures and arrange choices involving large, hypothetical “consequences,” whereas operant choice studies tend to arrange extended exposure to free-operant

contingencies involving real, though typically small, consequences. Disentangling all of these variables is beyond the scope of the present discussion, but it may be instructive to consider research areas in which some of the same variables operate. For example, delay-discounting studies, like cognitive-decision studies, typically involve brief procedures and large but hypothetical “consequences” (in choices between smaller-sooner and larger-later outcomes; see Green & Myerson, 2004). Some studies have found that money gains are discounted more heavily under delay than money losses (e.g., Chapman, 1996; Murphy, Vuchinich, & Simpson, 2001), suggesting disproportionate influence by the latter, although the effect has not always been replicated (e.g., Baker, Johnson, & Bickel, 2003; Odum, Madden, & Bickel, 2002).

Delay discounting also is invoked in human studies of self-control, which arrange for choices that are structured similarly to those of hypothetical delay-discounting procedures, although the consequences are real and contingency exposure can be substantial. Typically, humans make few impulsive (smaller-sooner) choices when positive reinforcement is involved, but may do so more readily when the outcomes involve avoiding aversive events (Navarick, 1982; Takahashi & Fujihara, 1995). This finding can be interpreted as showing that aversive events are *more* heavily discounted under conditions of delay, and therefore, in a sense, are *weaker* than appetitive events (although other interpretations are equally tenable; see Tice, Bratslavsky, & Baumeister, 2001).

Complicating matters further, the reinforcers in time-based intermittent reinforcement

schedules are defined in terms of both delay to and probability of occurrence. Recent evidence suggests that differential-impact effects are embedded in a complex interaction involving sign (gain versus loss), magnitude of outcome, and type of discounting (delay versus probability), such that gains are more steeply delay-discounted than losses only when outcomes are relatively small, and more steeply probability-discounted only when outcomes are relatively large (Estle, Green, Myerson, & Holt, 2006). Such findings point to a need to vary absolute reinforcer magnitude in studies like the present one to see if $\log b$ is systematically affected.

A Free-operant Differential-outcomes Effect?

The slopes of matching functions were consistently steeper under heterogeneous reinforcement than under homogeneous reinforcement (Figure 1 and 2; Table 3). Such an effect has uncertain bearing on differential-impact hypotheses but is anticipated by contingency-discriminability models of choice (e.g., Davison & Jenkins, 1985; Davison and Nevin, 1999; for a nontechnical introduction, see Magoon & Critchfield, 2006). These models represent what Davison and Nevin called a discriminative law of effect (DLOE), which assumes that undermatching is the norm in concurrent-schedule performance because the effects of reinforcer R_x on behavior B_x generalize partially to behavior B_y , and the effects of reinforcer R_y on behavior B_y generalize partially to behavior B_x . The amount of generalization depends on “the distinctiveness of the relation between behavior and reinforcement for one discriminated operant relative to another” and “reflects the joint effects of variables that influence response–reinforcer contingencies such as qualities or delays of the outcomes” (Davison & Nevin, p. 445).

When R_x and R_y are of the same type, as is typically the case in operant choice experiments, generalization should be more pronounced than when they are of different types. Davison and Nevin (1999) therefore proposed that matching involving heterogeneous reinforcers creates steeper slopes than matching involving homogeneous reinforcers. The present study may be the first to illustrate this predicted effect.

Davison and Nevin (1999) noted that DLOE models also anticipate the *differential-outcomes effect* that has been reported in conditional discrimination research (e.g., Trapold, 1970). In discrete-trials procedures, conditional discriminations tend to be acquired more efficiently, and to a higher level of accuracy, if the consequences associated with the stimulus–behavior relations are heterogeneous ($S_x: B_x \rightarrow R_x$ and $S_y: B_y \rightarrow R_y$) rather than homogeneous ($S_x: B_x \rightarrow R_x$ and $S_y: B_y \rightarrow R_x$). From the perspective of DLOE, this effect is a special case of a general contingency-discriminability rule. That is, by virtue of homogenous reinforcement, $B_x \rightarrow R_x$ and $B_y \rightarrow R_x$ are more similar to one another than are $B_x \rightarrow R_x$ and $B_y \rightarrow R_y$, leading to less generalization, and therefore sharper discriminations. In the latter case. Davison and Nevin were the first to specify this prediction, but in qualitative form it derives from the conceptual precepts of all DLOE models (Magoon & Critchfield, 2006). Previously, for reasons unrelated to DLOE, Goeters, Blakely, and Poling (1992) also proposed that a free-operant differential-outcomes effect is possible. Upon reviewing the literature, however, they found no reports of such an effect.

Viewing the present results as evidence of a free-operant differential outcomes effect supports not only the predictions of Davison and Nevin (1999) and Goeters et al. (1992) but also an interesting train of thought regarding theoretical views of negative reinforcement. Note that two-factor theories (e.g., Dinsmoor, 2001; Mowrer, 1947) view positive and negative reinforcement as different types of consequences, while one-factor theories (Herrnstein & Hineline, 1966; Rachlin & Herrnstein, 1969; Sidman, 1962) apparently do not. From a DLOE perspective, that slopes in the present study were elevated under heterogeneous reinforcement perhaps indicates that participants experienced positive and negative reinforcement as different types of consequences. If so, the finding of a free-operant differential-outcomes effect may be seen as broadly compatible with two-factor views of avoidance. It should be emphasized, however, that contingency-discriminability theorists have not linked their models to the theoretical debate over negative reinforcement; this speculation is strictly our own.

If participants behaved as if positive and negative reinforcement are different in type, it

remains unclear which aspects of these contingencies, as programmed in the present study, are functionally different. The two types of consequences, although programmed quite similarly, reflected at least three procedural asymmetries. First, slightly different instructions preceded initial exposure to the two types of consequences (see Appendix B). Second, the sign (+ or -) of the money outcome necessarily differed in the two cases. Third, the contingencies between responses and feedback messages associated with money outcomes differed in that feedback occurred frequently under positive reinforcement (when the contingency was met) and infrequently under negative reinforcement (when the contingency was not met; see "missed" reinforcers in Appendix A). The present experiment was not designed to disentangle the relative contributions of these factors.

Before the present findings can be confidently viewed as a free-operant differential outcomes effect, as per Davison and Nevin's (1999) account, several preliminary issues that are beyond the scope of the present investigation must be resolved. These are itemized below.

Baselines. The effects of heterogeneous reinforcement were evaluated via comparison to a homogeneous-reinforcement baseline function in which both concurrent schedules involved positive reinforcement. The present findings would be easier to interpret if more were known about the replicability of matching functions. Remarkably, given the long history of concurrent-schedules research, we know of no investigation in which test-retest reliability of steady-state matching functions was the primary focus. One general finding, that sensitivity (slope) tends to increase with schedule exposure (Todorov, Olivera Castro, Hannah, de Sa, & Barreto, 1983), can be ruled out as an explanation of the present study's slope effects because homogeneous-reinforcement and heterogeneous-reinforcement conditions alternated unsystematically for each participant. Nevertheless, any comparison of empirical matching functions depends on the assumption that each function is a reliable estimate of individual behavior tendencies.

In the present study, only the heterogeneous-reinforcement conditions involved negative reinforcement. Thus, slope effects might represent not a genuine differential-outcomes

effect, but rather something idiosyncratic about negative reinforcement in operant choice. In previous studies involving concurrent schedules of negative reinforcement (Baum, 1973; Ferrari & Todorov, 1980; Hutton, Gardner, & Lewis, 1978; Logue & de Villiers, 1978; Poling, 1978), sensitivity estimates were similar to what has been obtained from positive-reinforcement schedules, suggesting that the mere presence of negative reinforcement does not steepen slopes. Nevertheless, a productive next step would be to replicate the present study using baseline functions obtained from two concurrent negative-reinforcement schedules. Another strategy would be to employ two different types of positive reinforcement in heterogeneous reinforcement, although this approach creates challenges in equating the nominal value of qualitatively different reinforcers.

Discriminative stimuli. The present experiment was not designed with contingency-discriminability models in mind, and consequently it functions imperfectly as a test of the differential-outcomes prediction. The most thoroughly elaborated contingency-discriminability model (Davison & Nevin, 1999) assumes that the stimulus context in which concurrent schedules operate is constant across experimental conditions. Within conditions in the present study, the two component schedules were associated with targets of different colors, and these colors varied across conditions (see Appendix A). According to the Davison-Nevin model, which also assumes that discriminability of stimulus-behavior relations influences matching slopes, this might produce slope shifts if the discriminability of targets was systematically greater in heterogeneous-reinforcement conditions than in homogeneous-reinforcement conditions. Casual inspection of the target colors (Appendix A) suggests that this was not the case, but it would be prudent to replicate the study while holding stimulus features of the participant's display constant in order to rule out confounding of stimulus-behavior and behavior-reinforcer relations.

When $a > 1$. One hurdle to interpreting our results in terms of DLOE models is the fact that overmatching occurred in the heterogeneous-reinforcement condition for S447 and S448. For S448, slopes were significantly greater than 1 and for S447 the effect fell just

Table 5

Results of comparing empirical sensitivity (a) estimates to hypothetical functions with $a = 1$.

Participant	Responses			Time		
	F	DF	p	F	DF	p
<i>Equation 1</i>						
S447	6.79	1,4	.060	5.57	1,4	.078
S448	12.45	1,16	.003	5.23	1,16	.036
<i>Equation 4</i>						
S447	1.44	1,4	.297	0.88	1,13	.403
S448	3.31	1,16	.087	0.98	1,16	.336

Note. The comparison was not made for S449, for whom $a < 1$. Top: Estimates derived from fitting Equation 1 to heterogeneous-reinforcement data organized as negative reinforcement-schedule/positive-reinforcement schedule. See Figure 2 for parameter estimates. Bottom: Estimates derived from fitting Equation 4 to the same data, with punishment rates estimated from Appendix A (see text for details). See Table 6 for parameter estimates.

short of statistical significance (Table 5, top). In DLOE, strict matching ($a = 1$) is the theoretical limit of reinforcement sensitivity (Davison & Nevin, 1999). Aside from measurement error, Davison and Nevin cited two circumstances that could yield slopes greater than 1: changeover costs and punishment. In the former case, increases in changeover cost are known to steepen matching functions (Catania & Cutts, 1963), but in the present study changeover cost was constant under heterogeneous versus homogeneous reinforcement. Changeover costs might interact differently with negative reinforcement than with positive reinforcement, but we observed no systematic differences in changeover rates between heterogeneous and homogeneous conditions. In the latter case, although point losses originating in the negative reinforcement schedule might function as punishment, the mere presence of punishment does not assure slope changes; another possible effect is to create a bias away from the punished behavior (e.g., Critchfield, et al., 2003; McAdie, Foster, & Temple, 1996). It remains unclear how to reconcile steep slopes like those of S448 with the theoretical maximum assumed by DLOE models. As a starting point, it would help to evaluate the momentary dynamics of concurrent positive-negative reinforcement schedules in ways that the present study did not (a point to which we return shortly).

Reinforcer type versus amount. In considering the potential generality of a free-operant differential-outcomes effect, a useful point of comparison is provided by an investigation (McLean & Blampied, 2001) in which pigeons responded on concurrent schedules of rein-

forcement in both homogeneous conditions (reinforcer magnitude was equal for the two alternatives) and heterogeneous conditions (reinforcer magnitude was unequal for the two alternatives). Heterogeneous reinforcement produced no systematic change in slopes. Perhaps a free-operant differential-outcomes effect requires reinforcers to differ in type, not amount. DLOE models have not drawn this distinction explicitly, although Davison and Nevin (1999, Equation 20) speculated that between-schedules generalization based on reinforcer magnitude may be independent of that based on reinforcer frequency. By extension, generalization based on magnitude might be independent of that based on reinforcer type. Alternatively, magnitude effects may vary across studies because they interact with factors that are not systematically manipulated within studies. Already noted is the complex sign-magnitude-interaction of discounting research (Estle, et al., 2006). Recently Grace and Bragason found that effects of reinforcer magnitude on choice may depend on the temporal distribution of reinforcer delays in intermittent schedules (Grace & Bragason, 2005). Because of such effects, studies in which differential outcomes were defined via reinforcer magnitude may not provide the most instructive frame of reference for the present results.

Other Procedural Issues

The positive-reinforcement and negative-reinforcement contingencies of the present study were programmed similarly, but it remains possible that the results are idiosyncratically dependent on differences between

schedules that are conceptually unrelated to differential-impact or differential-outcomes effects. For example, consider that, exclusively in heterogeneous-reinforcement conditions, session earnings were supplemented to guard against very low earnings that might lead participants to quit the experiment. These supplements were not identified in the experimental instructions and did not appear on a participant's screen during experimental sessions. In pilot work (Critchfield & Magoon, 2001) we found that omitting the supplements did not eliminate slope effects similar to those described here. Nevertheless, the supplements altered session earnings and therefore could have influenced the economy of the experiment. The present study should be replicated without the money supplements.

Although, as noted in the Introduction, avoidance (unlike punishment) can operate without being imposed on other contingencies, it can interact with concurrently operating schedules. As is typical in the literature, in the present study positive reinforcement was programmed so that reinforcers could be obtained only by responding on a given schedule. What establishes avoidance as reinforcing, however, is the ongoing probability of an aversive event (e.g., Baron, 1991). To preserve this probability in the present study, point losses, if not prevented by responding, could occur at any time, including when participants were engaged with the concurrently available positive reinforcement schedule.

Our selection of a concurrent-schedules procedure may seem odd given that an alternative assay, the concurrent-chains procedure, is designed explicitly to prevent between-schedule interactions. The terminal-link schedules that determine preference operate in isolation and are entered only as a consequence of responding on competing, initial-link schedules. Preference is measured only in terms of initial-link responding (Mazur, 1991). It is easy to imagine a concurrent-chains version of the present study in which positive and negative reinforcement serve as the terminal links. Although events occurring in the terminal-link schedules would be fully independent, such a procedure is unsuitable for answering differential-impact questions about negative reinforcement because, in lay terms, concurrent-chains procedures may be said to assess the attractiveness of engaging in

terminal-link schedules. This is not necessarily synonymous with the strength of control over behavior that those schedules exert, once engaged. In concurrent-chains procedures, it seems unlikely that a participant ever would forego a money-gain, positive reinforcement terminal link to enter a money loss-avoidance schedule in which the best possible outcome (based on perfect avoidance) is a net monetary change of zero. Once engaged in such an avoidance schedule, however, the participant might respond quite vigorously. Thus, despite the fact that concurrent schedules partially intermingle the momentary dynamics of competing schedules, we employed them to require engagement with the negative-reinforcement schedule.

After employing a similar approach in their study of concurrent schedules of positive and negative reinforcement, Ruddle et al. (1982) estimated that most point losses associated with the avoidance schedule occurred while participants were engaged with the positive reinforcement schedule (because engaging the avoidance schedule tended to cancel point losses). If the same thing occurred in the present study, and these point losses adventitiously punished behavior maintained by positive reinforcement, then matching-function slopes might be increased (Davison & Nevin, 1999; Farley & Fantino, 1978)—an effect that clearly is relevant to the DLOE-inspired differential-outcomes hypothesis. Unfortunately, the present data, which depict the number of point losses ("missed" negative reinforcers in Appendix A) but not their timing, do not permit a direct evaluation of this possibility.

To illustrate the underlying issues, however, assume that *all* point losses served to punish responding on the positive-reinforcement schedule. Table 6 summarizes an analysis in which the data from heterogeneous-reinforcement conditions, organized as per Figure 2, were fitted to the following variant of the matching law in which reinforcement totals for each behavior option are decremented according to punishment frequency (e.g., de Villiers, 1980; Farley & Fantino, 1978; Davison & Nevin, 1999):

$$\log\left(\frac{R_x}{R_y}\right) = a \log\left(\frac{R_x - P_x}{R_y - P_y}\right) + \log b \quad (4)$$

Thus, point losses arising from the negative

Table 6

Results of fitting Equation 4 to heterogeneous-reinforcement data, organized as negative-reinforcement schedule/positive-reinforcement schedule.

Participant	Behavior measure	Fitted parameters			<i>a</i> vs. homogeneous		
		<i>a</i>	log <i>b</i>	%VAC	<i>F</i>	<i>DF</i>	<i>p</i>
S447	Responses	1.278	+1.108	93.8%	5.60	1,4	.077
	Time	1.142	−.087	96.5%	9.36	1,4	.038
S448	Responses	1.265	+1.126	90.4%	12.74	1,16	.003
	Time	1.131	+0.098	90.1%	12.14	1,16	.002
S449	Responses	0.631	−.014	90.9%	4.30	1,12	.058
	Time	0.497	−.070	90.1%	6.51	1,12	.025

Note. Also shown is a comparison of empirical sensitivity (*a*) estimates with those derived from fitting Equation 4 (which reduces to Equation 1 in the absence of punishment) to homogeneous-reinforcement data, organized as left-side schedule/right-side schedule. %VAC = percentage of variance in behavior allocation accounted for by Equation 4.

reinforcement schedule were assumed to reduce the value of the positive-reinforcement schedule ($P_x = 0$; P_y = the number of point losses resulting from imperfect avoidance). As Table 6 shows, the resulting slopes were less steep than in Figure 2, suggesting that adventitious punishment indeed could have contributed to slope effects. Table 2 (bottom) shows, however, that in four of six cases slopes remained significantly steeper than those for homogeneous reinforcement, with the remaining two cases narrowly missing significance. This suggests that slope effects of the present study may not be wholly artifactual. Finally, Table 5 (bottom) shows that none of the resulting log *b* estimates was significantly different from 0, which is broadly consistent with other analyses. Based on available data, it seems unlikely that adventitious punishment, if it occurred, dramatically affected the present results.

Unfortunately, the preceding exercise is hampered by a lack of contemporary empirical and conceptual guidance on how to integrate aversive events into the matching law. The past several decades of operant research have yielded many advances in the understanding of positive reinforcement, particularly as it influences choice, without much corresponding attention to aversive control (Critchfield & Rasmussen, 2007). Our selection of Equation 4 to represent punishment effects was somewhat arbitrary, as no consensus exists on the proper form of an operant choice model that incorporates punishment (Critchfield et al., 2003; Dinsmoor, 1998; Gray, Stafford, & Tallman, 1991; Lie & Alsop, 2007) or, for that matter, on the theoretical implications of any given model form (Dinsmoor, 1998). Additionally, because the present discussion focus-

es on how to interpret failures to cancel money losses under negative reinforcement, it is reasonable to ask about the interpretation of failures to earn money on a positive reinforcement schedule. These failures normally are ignored on the grounds that no feedback is associated with them, but this is conceptually inconsistent with the practice of treating unsignalled cancellations of money loss as functional events (i.e., negative reinforcers; de Villiers, 1972, 1974). Alternatively, failures to earn positive reinforcers might be considered aversive events (e.g., Lane & Cherek, 1999), although this is inconsistent with traditional analyses of positive reinforcement, which usually consider only obtained reinforcers (e.g., Mazur, 1991). The present study could be profitably extended by a replication that examined the momentary dynamics of behavior under concurrent schedules of positive and negative reinforcement, thereby allowing the precise temporal mapping of money losses that could have served as punishers. To be fully informative, however, such a study would require parallel development of operant choice models that describe how positive reinforcement combines with aversive events.

Conclusions

In recent decades basic research on aversive control has been undertaken infrequently in operant psychology (Baron, 1991; Critchfield & Rasmussen, 2007; Crosbie, 1998; Lerman & Vorndran, 2002). Perhaps as a result, many recent theories of the psychology of aversive control have evolved without reference to basic operant research (e.g., Frank, 2005;

Gehring & Willoughby, 2002; Gershoff, 2002; Holroyd & Coles, 2002; Kahneman & Tversky, 1979; Patterson, Kosson, & Newman, 1987). The differential-impact hypothesis that prompted the present investigation is a case in point. None of the recent major reviews examining this hypothesis (Baumeister et al., 2001; Rozin & Royzman, 2001; Taylor, 1991) cited operant studies of aversive control. Perhaps more important than the specifics of the present study's findings is the demonstration that interesting questions about aversive control of operant behavior remain to be adequately addressed. The fact that some of these questions have been better defined outside of operant psychology than within it can be viewed in two ways: as a failure of operant researchers to build upon past successes, or as an opportunity to demonstrate the contemporary relevance of operant aversive-control research. Additionally, as illustrated by the possible differential-outcomes effect described here, the process of addressing questions that matter to scholars outside of operant psychology may sometimes lead back to questions of interest to operant theory.

REFERENCES

- Baker, F., Johnson, M. W., & Bickel, W. K. (2003). Delay discounting in current and never-before cigarette smokers: Similarities and differences across commodity, sign, and magnitude. *Journal of Abnormal Psychology*, 112, 382–392.
- Baron, A. (1991). Avoidance and punishment. In I. H. Iversen & K. A. Lattal (Eds.), *Experimental analysis of behavior, Part 1* (pp. 173–217). Amsterdam: Elsevier.
- Baum, W. M. (1973). Time allocation and negative reinforcement. *Journal of the Experimental Analysis of Behavior*, 20, 313–322.
- Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, 22, 231–242.
- Baum, W. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior*, 32, 269–281.
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, 5, 323–370.
- Catania, A. C., & Cutts, D. (1963). Experimental control of superstitious responding in humans. *Journal of the Experimental Analysis of Behavior*, 6, 203–208.
- Chao, E. T. (1984). Hedonic scaling of sugars using concurrent operant schedules with dogs. *Neuroscience & Biobehavioral Reviews*, 8, 225–229.
- Chapman, G. B. (1996). Temporal discounting and utility for health and money. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 22, 771–791.
- Christensen, J., Parker, S., Silberberg, A., & Hursh, S. (1998). Trade-offs in choice between risk and delay depend on monetary amounts. *Journal of the Experimental Analysis of Behavior*, 69, 123–139.
- Costantini, A. F., & Hoving, K. L. (1973). The effectiveness of reward and punishment contingencies on response inhibition. *Journal of Experimental Child Psychology*, 16, 484–494.
- Critchfield, T. S., & Magoon, M. A. (2001). On the differential impact of positive and negative reinforcement. *Experimental Analysis of Human Behavior Bulletin*, 19, 16–18.
- Critchfield, T. S., Paletz, E. M., MacAleese, K. R., & Newland, M. C. (2003). Punishment in human choice: Direct or competitive suppression? *Journal of the Experimental Analysis of Behavior*, 80, 1–27.
- Critchfield, T. S., & Rasmussen, E. R. (2007). It's aversive to have an incomplete science of behavior. *Mexican Journal of Behavior Analysis*, 33, 1–6.
- Critchfield, T. S., Schlund, M., & Ecott, C. L. (2000). A procedure for using bonus course credit to establish points as reinforcers for human subjects. *Experimental Analysis of Human Behavior Bulletin*, 18, 15–18.
- Crosbie, J. (1998). Negative reinforcement and punishment. In K. A. Lattal & M. Perone (Eds.), *Handbook of research methods in human operant behavior* (pp. 163–198). New York: Plenum.
- Daniels, A. C. (1994). *Bringing out the best in people: How to apply the astonishing power of positive reinforcement*. New York: McGraw-Hill.
- Davison, M. C., & Jenkins, P. E. (1985). Stimulus discriminability, contingency discriminability, and schedule performance. *Animal Learning & Behavior*, 13, 77–84.
- Davison, M., & McCarthy, D. (1988). *The matching law: A research review*. Hillsdale, NJ: Erlbaum.
- Davison, M., & Nevin, J. A. (1999). Stimuli, reinforcers, and behavior: An integration. *Journal of the Experimental Analysis of Behavior*, 71, 439–482.
- deVilliers, P. A. (1972). Reinforcement and response rate interaction in multiple random-interval avoidance schedules. *Journal of the Experimental Analysis of Behavior*, 18, 499–507.
- deVilliers, P. A. (1974). The law of effect and avoidance: A quantitative relationship between response rate and shock-frequency reduction. *Journal of the Experimental Analysis of Behavior*, 21, 223–235.
- de Villiers, P. A. (1980). Toward a quantitative theory of punishment. *Journal of the Experimental Analysis of Behavior*, 33, 15–25.
- Dinsmoor, J. (1954). Punishment I: The avoidance hypothesis. *Psychological Review*, 61, 34–46.
- Dinsmoor, J. A. (1998). Punishment. In W. T. O'Donohue (Ed.), *Learning and behavior therapy* (pp. 188–204). Needham Heights, MA: Allyn & Bacon.
- Dinsmoor, J. A. (2001). Still no evidence for temporally extended shock-frequency reduction as a reinforcer. *Journal of the Experimental Analysis of Behavior*, 75, 367–378.
- Estle, S. J., Green, L., Myerson, J., & Holt, D. D. (2006). Differential effects of amount on temporal and probability discounting of gains and losses. *Memory & Cognition*, 34, 914–928.
- Farley, J., & Fantino, E. (1978). The symmetrical law of effect and the matching relation in choice behavior. *Journal of the Experimental Analysis of Behavior*, 29, 37–60.

- Ferrari, E. A., & Todorov, J. C. (1980). Concurrent avoidance of shocks by pigeons pecking a key. *Journal of the Experimental Analysis of Behavior*, 34, 329–333.
- Fleshler, M., & Hoffman, H. S. (1962). A progression for generating variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 5, 529–530.
- Frank, M. J. (2005). Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and nonmedicated Parkinsonism. *Journal of Cognitive Neuroscience*, 17, 51–72.
- Gehring, W. J., & Willoughby, A. R. (2002). The medial frontal cortex and the rapid processing of monetary gains and losses. *Science*, 295, 2279–2282.
- Gershoff, E. T. (2002). Corporal punishment by parents and associated child behaviors and experiences: A meta-analytic and theoretical review. *Psychological Bulletin*, 128, 539–579.
- Goeters, S., Blakely, E., & Poling, A. (1992). The differential outcomes effect. *The Psychological Record*, 42, 389–411.
- Grace, R. C., & Bragason, O. (2005). Does sensitivity to magnitude depend on the temporal distribution of reinforcement? *Journal of the Experimental Analysis of Behavior*, 83, 169–184.
- Gray, L. N., Stafford, M. C., & Tallman, I. (1991). Rewards and punishment in complex human choices. *Social Psychology Quarterly*, 54, 318–329.
- Green, L., & Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin*, 130, 769–792.
- Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, 13, 243–266.
- Herrnstein, R. J., & Hineline, P. N. (1966). Negative reinforcement as shock-frequency reduction. *Journal of the Experimental Analysis of Behavior*, 9, 421–430.
- Hineline, P. N. (1984). Aversive control: A separate domain? *Journal of the Experimental Analysis of Behavior*, 42, 495–509.
- Hollard, V., & Davison, M. C. (1971). Preference for qualitatively different reinforcers. *Journal of the Experimental Analysis of Behavior*, 16, 375–380.
- Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and error-related negativity. *Psychological Review*, 109, 679–709.
- Hutton, L., Gardner, E. T., & Lewis, P. (1978). Matching with a key-peck response in concurrent negative reinforcement schedules. *Journal of the Experimental Analysis of Behavior*, 30, 225–230.
- Kahneman, D., Knetch, J. L., & Thaler, R. H. (1990). Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy*, 98, 1325–1348.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–291.
- Kollins, S. H., Newland, M. C., & Critchfield, T. S. (1997). Human sensitivity to reinforcement in operant choice: How much do consequences matter? *Psychonomic Bulletin & Review*, 4, 208–220. Erratum: *Psychonomic Bulletin & Review*, 4, 431.
- Lane, S. D., & Cherek, D. R. (1999). Decision under conditions of risk: Exploring some parameters and quantitative models. *Experimental Analysis of Human Behavior Bulletin*, 17, 15–19.
- Lerman, D. C., & Vorndran, C. M. (2002). On the status of knowledge for using punishment: Implications for treating behavior disorders. *Journal of Applied Behavior Analysis*, 35, 431–464.
- Lie, C., & Alsop, B. (2007). Reinforcement and punishment in behavioral models of signal detection. *Mexican Journal of Behavior Analysis*, 33, 45–55.
- Logue, A. W., & de Villiers, P. A. (1978). Matching in concurrent variable-interval avoidance schedules. *Journal of the Experimental Analysis of Behavior*, 29, 61–66.
- Logue, A. W., & de Villiers, P. A. (1981). Matching of behavior maintained by concurrent shock avoidance and food reinforcement. *Behaviour Analysis Letters*, 1, 247–258.
- Madden, G. J., & Perone, M. (1999). Human sensitivity to concurrent schedules of reinforcement: Effects of observing schedule-correlated stimuli. *Journal of the Experimental Analysis of Behavior*, 71, 303–318.
- Madden, G. J., & Perone, M. (2003). Effects of alternative reinforcement on human behavior: The source does matter. *Journal of the Experimental Analysis of Behavior*, 79, 193–206.
- Magoon, M. A., & Critchfield, T. S. (2006). The Discriminative Law of Effect: Rethinking the concept of contingency. *European Journal of Behavior Analysis*, 7, 159–171.
- Matthews, L. R., & Temple, W. (1979). Concurrent schedule assessment of food preference in cows. *Journal of the Experimental Analysis of Behavior*, 32, 245–254.
- Mazur, J. (1991). Choice. In I. H. Iversen & K. A. Lattal (Eds.), *Experimental analysis of behavior, Part 1* (pp. 219–250). New York: Elsevier.
- McAdie, T. M., Foster, T. M., & Temple, W. (1996). Concurrent schedules: Quantifying the aversiveness of noise. *Journal of the Experimental Analysis of Behavior*, 65, 37–55.
- McLean, A. P., & Blampied, N. M. (2001). Sensitivity to relative reinforcer rate in concurrent schedules: Independence from relative and absolute reinforcer duration. *Journal of the Experimental Analysis of Behavior*, 75, 25–42.
- Meyer, W. J., & Offenhach, S. (1962). Effectiveness of reward and punishment as a function of task complexity. *Journal of Comparative and Physiological Psychology*, 55, 532–534.
- Miller, H. L., Jr. (1976). Matching-based hedonic scaling in the pigeon. *Journal of the Experimental Analysis of Behavior*, 26, 335–347.
- Motoulsky, H., & Christopoulos, A. (2006). *Fitting models to biological data using linear and nonlinear regression: A practical guide to curve fitting*. San Diego: GraphPad Software.
- Mowrer, O. H. (1947). On the dual nature of learning: A re-interpretation of “conditioning” and “problem solving.” *Harvard Educational Review*, 17, 102–148.
- Murphy, J. G., Vuchinich, R. E., & Simpson, C. A. (2001). Delayed reward and cost discounting. *The Psychological Record*, 51, 571–588.
- Navarick, D. J. (1982). Negative reinforcement and choice in humans. *Learning and Motivation*, 13, 361–377.
- Odum, A. L., Madden, G. J., & Bickel, W. K. (2002). Discounting of delayed health gains and losses by current, never- and ex-smokers of cigarettes. *Nicotine & Tobacco Research*, 4, 295–303.

- Patterson, C. M., Kosson, D. S., & Newman, J. P. (1987). Reaction to punishment, reflectivity, and passive avoidance learning in extraverts. *Journal of Personality and Social Psychology*, 52, 565–575.
- Penney, R. K. (1968). Effects of reward and punishment on children's orientation and discrimination learning. *Journal of Experimental Psychology*, 75, 140–142.
- Penney, R. K., & Lupton, A. A. (1961). Children's discrimination learning as a function of reward and punishment. *Journal of Comparative and Physiological Psychology*, 54, 449–451.
- Pilgrim, C. (1988). The human subject. In K. A. Lattal & M. Perone (Eds.), *Handbook of research methods in human operant behavior* (pp. 15–44). New York: Plenum.
- Poling, A. (1978). Performance of rats under concurrent variable-interval schedules of negative reinforcement. *Journal of the Experimental Analysis of Behavior*, 30, 31–36.
- Rachlin, H., & Herrnstein, R. J. (1969). Hedonism revisited: On the negative law of effect. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 83–109). New York: Appleton-Century-Crofts.
- Rachlin, H. (1989). *Judgment, decision, and choice*. New York: Freeman.
- Rachlin, H., Logue, A. W., Gibbon, J., & Frankel, M. (1986). Cognition and behavior in studies of choice. *Psychological Review*, 93, 33–45.
- Robinson, J. K. (1992). Quantitative analyses of choice in the rat and pigeon. *The Psychological Record*, 42, 437–445.
- Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review*, 5, 296–320.
- Ruddle, H. V., Bradshaw, C. M., & Szabadi, E. (1981). Performance of humans in variable-interval avoidance schedules programmed singly, and concurrently with variable-interval schedules of positive reinforcement. *Quarterly Journal of Experimental Psychology*, 33, 213–226.
- Ruddle, H. V., Bradshaw, C. M., Szabadi, E., & Foster, T. M. (1982). Performance of humans in concurrent avoidance/positive-reinforcement schedules. *Journal of the Experimental Analysis of Behavior*, 38, 51–61.
- Sidman, M. (1962). Reduction of shock frequency as reinforcement for avoidance behavior. *Journal of the Experimental Analysis of Behavior*, 5, 247–257.
- Sidman, M. (1966). Avoidance behavior. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 448–498). New York: Appleton-Century-Crofts.
- Sidman, M. (1989). *Coercion and its fallout*. Boston: Authors Cooperative.
- Skinner, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Croft.
- Skinner, B. F. (1953). *Science and human behavior*. New York: Free Press.
- Spence, J. T. (1966). Verbal discrimination performance as a function of instruction and verbal reinforcement combination in normal and retarded children. *Child Development*, 37, 269–281.
- Takahashi, M., & Fujihara, T. (1995). Self-control and choice in humans: Effects of type, amount, and delay of reinforcers. *Learning and Motivation*, 26, 183–202.
- Taylor, S. E. (1991). Asymmetrical effects of positive and negative events: The mobilization-minimization hypothesis. *Psychological Bulletin*, 110, 67–85.
- Thorndike, E. L. (1911). *Animal intelligence*. New York: Macmillan.
- Tice, D. M., Bratslavsky, E., & Baumeister, R. F. (2001). Emotional distress regulation takes precedence over impulse control: If you feel bad, do it! *Journal of Personality & Social Psychology*, 80, 53–67.
- Tindall, R. C., & Ratliff, R. G. (1974). Interaction of reinforcement conditions and developmental level in a two-choice discrimination task with children. *Journal of Experimental Child Psychology*, 18, 183–189.
- Todorov, J. C., Olivera Castro, J. M., Hanna, E. S., de Sa, M. C. N. B., & Barreto, M. (1983). Choice, experience, and the generalized matching law. *Journal of the Experimental Analysis of Behavior*, 40, 99–111.
- Trapold, M. A. (1970). Are expectancies based on different positive reinforcing events discriminably different? *Learning and Motivation*, 1, 129–140.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211, 453–458.
- Zar, J. H. (1999). *Biostatistical analysis* (4th ed.). Upper Saddle River, NJ: Prentice Hall.

Received: March 14, 2005

Final Acceptance: February 25, 2008

APPENDIX A

Experimental Conditions and Terminal Data (means from the last four sessions per condition). Conditions: For each participant, conditions are listed in order of relative richness of the schedule with the higher programmed reinforcement rate (where a difference existed). The designations x and y refer to the schedules represented in Equation 1. Schedules: N = negative reinforcement and P = positive reinforcement. Numerals represent the approximate programmed reinforcement ratio. See Table 1 for details on schedule values. Colors: Four shades (ranging from 1 = lightest to 4 = darkest) of blue (B), green (G), yellow (Y), and violet (V) were used, for a total of 16 unique colors. Sides: L = left side of participant's screen, R = right side of participant's screen. Money Supplement: Amount added, per session, to participant's overall total. Sessions: Number required to meet the stability criterion. Obtained Reinforcers: Those contingent upon a response occurring at any time during a targeted interval of the variable-cycle schedule. Under negative reinforcement, this produced unsignaled cancellation of a scheduled money loss. Under positive reinforcement, this produced a signaled money gain. Missed Reinforcers: Those programmed under variable-cycle schedules but not obtained because no response occurred during the targeted interval. Under negative reinforcement, this produced a signaled money loss. Under positive reinforcement, this produced an unsignaled loss of a scheduled opportunity to earn money.

Participant	Condition	Schedules $x:y$	Colors $x:y$	Sides $x:y$	Money suppl- ement	Sessions	Reinforcers per min				Behavior allocation per min			
							Obtained		Missed		Responses		Time (s)	
							x	y	x	y	x	y	x	y
S447	4	N9:P1	G1:V3	R:L	\$0.82	8	5.53	0.47	0.07	0.07	141.13	2.43	57.23	2.69
	5	P9:P1	G4:V1	L:R	—	4	5.05	0.45	0.38	0.08	130.35	18.93	50.83	9.08
	8	N4:P1	P4:G2	L:R	\$0.69	6	4.30	0.83	0.38	0.35	105.95	15.53	50.67	9.17
	7	P4:P1	Y4:B4	L:R	—	5	3.60	0.93	0.90	0.28	81.38	51.08	38.14	21.69
	2	N1:P1	V3:G3	L:R	\$0.45	5	2.53	2.28	0.48	0.70	75.05	43.43	37.36	22.48
	1	P1:P1	B3:Y3	L:R	—	4	2.13	2.03	0.88	0.95	62.75	54.58	30.73	29.09
	6	N1:P9	B2:Y2	L:R	\$0.09	8	0.35	5.00	0.15	0.45	3.95	110.80	2.11	57.67
S448	3	P1:P9	G2:B1	L:R	—	4	0.50	4.50	0.03	0.90	33.60	109.83	14.66	45.21
	8	N9:P1	B4:V1	L:R	\$0.82	4	5.18	0.50	0.35	0.05	49.15	2.95	55.22	4.57
	18	P9:P1	G3:V3	L:R	—	6	4.65	0.50	0.90	0.03	41.48	15.35	42.80	17.12
	13	N6:P1	V1:B3	R:L	\$0.75	4	4.63	0.78	0.25	0.13	48.45	3.05	54.65	5.29
	9	P6:P1	B2:G4	L:R	—	4	3.80	0.85	1.15	0.08	36.18	17.05	39.32	20.62
	4	N4:P1	B4:Y4	R:L	\$0.69	6	4.33	0.80	0.20	0.38	46.35	4.10	53.09	6.81
	3	P4:P1	B2:Y2	L:R	—	6	3.48	0.93	1.15	0.28	32.50	21.38	35.47	24.40
	17	N2:P1	Y3:B2	L:R	\$0.60	7	3.75	1.40	0.13	0.48	46.00	7.20	51.35	8.57
	15	P2:P1	B3:G4	L:R	—	7	3.15	1.80	0.78	0.13	29.25	27.15	31.10	28.84
	6	P1:P1	G3:P4	L:R	—	4	2.58	2.50	0.38	0.48	22.30	21.88	30.27	29.66
	16	N1:P1	V1:G1	R:L	\$0.45	8	2.83	1.90	0.18	1.08	42.25	5.75	49.93	10.00
	19	N1:P2	Y1:B1	R:L	\$0.30	8	1.45	3.25	0.33	0.70	21.13	34.30	23.90	35.99
	12	P2:P3	G4:Y4	L:R	—	5	2.00	2.50	0.35	0.85	28.28	25.58	31.76	28.18
	20	N2:P3	B4:V1	R:L	\$0.38	6	2.25	2.68	0.13	1.13	35.38	16.45	41.07	18.92
	14	P1:P2	V3:G2	L:R	—	4	1.68	3.10	0.15	0.83	25.40	27.65	29.28	30.65
	2	N1:P4	B1:G1	L:R	\$0.18	7	0.10	4.53	1.00	0.00	0.70	57.28	1.02	58.88
	1	P1:P4	G3:Y4	L:R	—	7	0.98	3.45	0.23	1.08	23.28	27.83	28.60	31.26
	10	N1:P6	B1:Y3	L:R	\$0.12	7	0.15	4.88	0.70	0.10	0.25	57.98	0.41	59.50
	11	P1:P6	V2:Y1	L:R	—	5	0.83	3.93	0.10	1.03	20.05	32.88	23.76	36.18
	5	N1:P9	B2:V3	R:L	\$0.09	11	0.53	5.18	0.03	0.30	4.40	51.18	6.49	53.43
S449	7	P1:P9	Y1:B1	L:R	—	12	0.45	5.10	0.10	0.40	5.45	49.30	8.01	51.91
	15	N9:P1	B4:V1	R:L	\$0.82	8	4.78	0.58	0.73	0.03	7.23	45.03	10.13	49.79
	5	P9:P1	V1:G2	L:R	—	4	4.45	0.58	0.85	0.00	33.33	11.13	40.63	19.19
	12	N6:P1	B4:Y4	L:R	\$0.75	7	4.50	0.68	0.50	0.03	42.28	9.23	47.70	12.23
	4	P6:P1	V2:B3	L:R	—	4	4.53	0.83	0.48	0.08	33.35	11.30	40.35	19.54
	14	N4:P1	G4:B2	L:R	\$0.69	6	4.18	0.83	0.38	0.38	41.78	8.53	47.19	12.73
	1	P4:P1	Y4:B2	L:R	—	4	3.50	1.00	1.03	0.18	31.03	18.93	35.89	24.01
	13	N2:P1	V1:G3	L:R	\$0.60	5	3.28	1.40	0.73	0.50	34.60	15.85	39.60	20.33
	11	P2:P1	G4:B2	L:R	—	8	3.65	1.60	0.28	0.25	23.35	16.83	33.11	26.80
	6	P1:P1	B2:Y4	L:R	—	7	2.65	2.63	0.30	0.35	18.80	18.30	30.37	29.52
	16	N1:P1	Y3:B2	R:L	\$0.45	4	2.58	2.63	0.40	0.35	20.23	19.15	29.64	30.29
	3	N1:P2	V1:B3	R:L	\$0.30	4	1.73	3.50	0.20	0.48	11.18	32.98	19.01	40.89
	7	P1:P2	Y1:V3	L:R	—	5	1.68	3.40	0.23	0.53	16.23	25.73	25.30	34.62
	8	N1:P6	Y2:B1	R:L	\$0.12	5	0.93	4.00	0.00	0.98	12.93	25.03	24.69	35.19
	10	P1:P6	G2:Y3	L:R	—	6	0.90	3.95	0.00	0.98	9.73	14.98	25.93	33.98
	2	N1:P9	G3:B1	R:L	\$0.09	5	0.53	4.70	0.03	0.65	8.30	37.25	15.44	44.46
	9	P1:P9	B4:Y3	L:R	—	4	0.58	4.40	0.00	1.10	11.60	16.95	25.62	34.24

APPENDIX B

PRELIMINARY TRAINING AND MAIN-EXPERIMENT
INSTRUCTIONS

Preliminary training began with a 5-min session during which a VC 12-s schedule of positive reinforcement operated on the right or left work area only. No target was displayed in the other work area, and no CO button was available. Participants read these instructions:

Sometimes clicking your mouse can earn you money. The upcoming session provides an example. If you should earn money, a flashing message will show this on your screen.

A subsequent 5-min session involving a single VC 12-s schedule of negative reinforcement, operating in the previously unused work area, was preceded by these instructions:

Sometimes clicking your mouse can keep you from losing money. The upcoming session provides an example. If you should lose money, a flashing message will show this on your screen.

A final set of written instructions was provided prior to the first concurrent-schedules session.

From now on, both sides of your screen will be active. Note that clicking the two sides may affect your earnings differently. Each side can either earn you money, or keep you from losing money. You can respond as much or as little as you like on either side, and use whatever strategy you like overall. It is up to you to figure out how to work each side to your best advantage.

After informed consent was obtained and preliminary training completed, participants completed a 2-hr phase designed to screen for sensitivity to differential reinforcement rates in concurrent schedules. Because the experimental design required the comparison of matching functions across two sets of conditions, it was important to identify volunteers who generally allocated more behavior to the richer of two schedules (which is not always a given with human participants; see Kollins, Newland, & Critchfield, 1997). During the screening phase, concurrent VC 12 s VC 60 s schedules of positive reinforcement (5:1 programmed ratio) operated in the work areas and produced points exchangeable for course

extra credit (see Critchfield, Schlund, & Ecott, 2000) rather than money. The value of the credit depended on course contingencies set by individual instructors and thus could vary across participants. Initially, the richer schedule was arbitrarily assigned to one of the work areas. Once visual inspection suggested a clear preference for this schedule, the schedules associated with the two work areas were reversed. The research protocol called for exclusion of any individual who failed to show a clear preference or to reverse this preference in accordance with schedule reversals, but based on this minimal screening all volunteers appeared to be sensitive to reinforcement-rate differentials.

Prior to the start of the first session of the main experiment, participants read the following instructions:

After you click the <Ready> button on the first screen you see, your "work" screen will appear. On this "work" screen you will notice boxes moving around on two different sides of the screen. Your job is to use the mouse to click on these moving boxes. Once you start working on one side, the other side of the screen will turn black. This means that your mouse is turned on for the white side, and turned off for the black side. Either side can be black or white, depending on where you are working at the moment. Clicking on the white side will help you earn money, or avoid losing money. Clicking on the black side will not help you, because the mouse is turned off for that side. While a side is black, you cannot earn money that may be available there. Similarly, while a side is black, you cannot avoid money losses that happen there.

To change a side from black to white, you will need to click several times on the small "Change" box in the middle of your screen. When you have clicked this box enough times, the black side will turn white, and the mouse will be turned on there. The mouse will be off for the other side.

Thus, the mouse is always on for only one side at a time. You can change sides whenever you want by clicking on the "Change" box. There are no rules for how you should respond except that you want to maximize the number of points you get as quickly as you can. An experimenter will record how much you earn at the end of each session. Please do not hesitate to ask the experimenter any questions you may have.